

Network Delay Inference from Additive Metrics

Shankar Bhamidi

Department of Statistics and OR
UNC Chapel Hill

Ram Rajagopal

Department of EECS
UC Berkeley

Sébastien Roch

Department of Mathematics
UC Los Angeles

July 28, 2009

Abstract

We use computational phylogenetic techniques to solve a central problem in inferential network monitoring. More precisely, we design a novel algorithm for multicast-based delay inference, that is, the problem of reconstructing delay characteristics of a network from end-to-end delay measurements on network paths. Our inference algorithm is based on additive metric techniques used in phylogenetics. It runs in polynomial time and requires a sample of size only $\text{poly}(\log n)$. We also show how to recover the topology of the routing tree.

1 Introduction

Network tomography. Inferential network monitoring—also known as network tomography [27]—consists in reconstructing various properties of large communication networks from indirect measurements in order to facilitate the management of these networks. Network inference can be achieved by two general approaches. In the *internal* approach, one takes measurements directly at the edges and nodes of the network. This approach suffers from several drawbacks: the network operator may not allow access to internal devices of the network or may not make public measurements on them; the routers may not have the technological capabilities to perform the required measurements; direct measurements may create extra computational burden as well as congestion in the network. This has led some in the networking community to consider instead the *external* approach. In this case, one uses so-called “end-to-end” measurements, e.g., measurements of delays or rate of packet drops between nodes in the network, and seeks to infer the desired network properties from them. This gives rise to an inverse problem similar to tomographic image reconstruction.

Our aim in this paper is to propose a novel approach to this problem. We focus on multicast-based inference. Multicast routing consists in sending a packet from a source to a set of receivers through a routing tree. The packet is duplicated at each branch point and sent further down the tree. The routing tree is generally unknown to the user. The idea is to use inherent correlation of measurements between different receivers to reconstruct the topology of the routing tree as well as to estimate link properties of this tree. The main link property we consider here is the delay distribution. The multicast inference approach was introduced in [4, 23].

A core difficulty of the problem is to devise efficient, scalable algorithms which consistently estimate the desired network properties. Several techniques have been used in the network tomography literature, notably maximum pseudo-likelihood, EM algorithms and Markov chain Monte Carlo methods. See [6] for a detailed survey and bibliographic references. In this paper, we introduce a new methodology for multicast delay inference inspired by techniques from the field of *phylogenetics* in biology, that is, the reconstruction of evolutionary trees from molecular data. Our methodology has the advantage of being provably consistent and computationally efficient. It also uses a small asymptotic sample size. This is crucial to reduce the burden on the network as well as to obtain a consistent “snapshot” of the network, which is intrinsically dynamic in nature. Typical networks undergo sporadic medium

to large-scale changes in structure over time, therefore algorithms with low sample complexity are essential. Concurrently to our work, Liang et al. [15] used similar ideas to tackle the related multicast packet loss inference problem. Also, Ni and Tatikonda [19] independently proposed a Markov-based inference algorithm similar to ours for multicast delay inference—although our work appears to be the first rigorous analysis of the sample complexity of this approach. See Section 1.2 for a precise statement of our results and Section 1.3 for a discussion of previous work. The results detailed here were first announced in [2].

Phylogeny background. A core problem in evolutionary biology is the inference of evolutionary histories of organisms from molecular data. Evolution is usually represented by a tree where branching points indicate speciation events. The root of the tree is the common ancestor to all species in the tree and the leaves are contemporary (extant) species. Molecular data is assumed to evolve according to a standard Markov model. The *phylogenetic reconstruction* problem is the following. From measurement of sequences of molecular data at the leaves, one seeks to reconstruct the topology of the evolutionary tree as well as mutation characteristics along the branches. See [11] and [25] for an overview of the field of phylogenetics.

Various statistical and computational techniques have been used to solve the phylogenetic reconstruction problem: maximum likelihood, bayesian, parsimony, and distance-based methods. In this paper, we adapt and extend *distance-based techniques* to deal with a class of models introduced in [23] in connection with the multicast network inference problem—this new class of models is similar to the Markov models used in phylogenetics but presents challenges of its own. The main idea in distance-based methods is to define a so-called tree metric from mutation parameters. A tree metric is a metric on the leaves of the tree which can be realized as a path metric on a corresponding weighted tree. (See Section 2 for more details.) After being estimated, the metric allows the reconstruction of the tree and its characteristics. A main advantage of this approach is that it leads to computationally efficient algorithms with provable sample requirement guarantees.

1.1 Basic Definitions

A broadcasting process on a tree. We now give a more formal statement of the multicast inference problem introduced in [23]. Let $T = (V, E)$ be a tree on $n + 1$ leaves L —representing the *routing tree*—and let $\{d_e\}_{e \in E}$ be a set of independent positive random variables on the edges—representing the *delays*. Leaf 0, the *source*, is the root of the tree. The remaining n leaves are the *receivers*. We assume that all internal nodes have degree at least 3.

A realization of the *multicast delay process* works as follows: the root sends a packet to the receivers through the routing tree; at every branching point, the packet is duplicated; on every link e , an independent random delay d_e is experienced by the packet. More formally, we define the multicast delay process $\{D_u\}_{u \in V}$ as follows. Let P_{ij} be the path (set of edges) between nodes i and j in T . For a node u , let

$$D_u = \sum_{e \in P_{0u}} d_e. \quad (1)$$

Note that D_u is the total delay at node u in the network.

The multicast inference problem. The tree and delay distributions are actually unknown to us. We are only given access to k independent samples of delays at the leaves $\{D_a^1\}_{a \in L}, \dots, \{D_a^k\}_{a \in L}$. Our goal is to reconstruct the routing tree and estimate the delay distributions using these samples. We now define more precisely what we mean by the estimation of the delay distributions. In this work, we assume that each edge delay distribution (in general, different) is characterized by a constant number, say $J - 1 > 1$ (independent of n), of consecutive central moments. That is, we assume there are

characteristic moments

$$w_e^{(j)} = \mathbb{E} \left[(d_e - \mathbb{E}[d_e])^j \right],$$

for all $e \in E$ and $2 \leq j \leq J$. Our goal is to estimate these moments within a fixed accuracy. More formally, we make the following assumption. We first need a definition.

Definition 1 (Regular Families) Let $\varepsilon > 0$ and $J \geq 2$ be fixed. Let $\mathcal{Q} = \{Q_\theta\}_{\theta \in \Theta}$ be a family of distributions on \mathbb{R} parametrized by $\theta \in \Theta$ where Θ is a subset of an Euclidean space. Let $\{w^{(j)}(\theta)\}_{\{2 \leq j \leq J\}}$ be the first $J - 1$ central moments of Q_θ . We say that the family \mathcal{Q} is (ε, J) -regular if there exists a map Ψ from \mathbb{R}^{J-1} to Θ and a $\delta > 0$ such that if the vector $\hat{\mathbf{w}} = \{\hat{w}^{(j)}\}_{\{2 \leq j \leq J\}}$ satisfies

$$\left| \hat{w}^{(j)} - w^{(j)}(\theta) \right| \leq \delta$$

for all $2 \leq j \leq J$, then

$$\|Q_\theta - Q_{\Psi(\hat{\mathbf{w}})}\|_1 \leq \varepsilon.$$

In Appendix A, we give simple examples of regular families.

Assumption 1 (Regularity and Boundedness) Let $\varepsilon > 0$ and $J \geq 2$ be fixed (independent of n). We assume that all edge delay distributions are from a fixed (ε, J) -regular family of distributions. Furthermore, we assume that the delays are uniformly bounded, namely there is a constant $M > 0$ independent of n such that for all $e \in E$, $d_e \in [0, M]$.

This framework is simple enough to be tractable yet general enough to accommodate large classes of distributions: parametrized distributions, e.g., beta distributions; and nonparametrized distributions, e.g., discretized distributions on $\{0, 1, \dots, M\}$. Further we need the following assumption.

Assumption 2 (Lower Bound on Second Moment) We assume that there is a constant $f > 0$ (independent of n) such that for all $e \in E$,

$$w_e^{(2)} \geq f.$$

To sum up, the *multicast inference problem* is defined as follows.

Definition 2 (Multicast Inference Problem, Moment Version) Let $\varepsilon > 0$ and $J \geq 2$ be fixed. The multicast inference problem consists in the following. Let T and $\{w_e^{(j)}\}_{\{e \in E, 2 \leq j \leq J\}}$ be any tree (with internal degrees at least 3) and set of central moments on edges. Given samples of delays at the leaves, we are required to:

1. **Tree Reconstruction.** Recover T .
2. **Moment Estimation.** Estimate all characteristic moments $\{w_e^{(j)}\}_{\{e \in E, 2 \leq j \leq J\}}$ within ε .

Remark 1 As noted by Lo Presti et al. [23], the means of the edge delay distributions are, in general, unidentifiable. See Figure 1 for an illustration. In particular, one cannot hope to recover the deterministic transmission delay on each link. But, as noted in [23], this is not a major issue. Indeed, in practice, one is only interested in the variable portion of the delay, that is, the portion resulting from traffic. To restore identifiability, Lo Presti et al. proceed by subtracting the lowest observed delay on each receiver, in order to remove the (estimated) deterministic component of the delay. They further assume that the variable portion of the delay “starts at 0.” We also make this last assumption (see our examples of regular delay distributions in Appendix A). However, instead of subtracting the minimum observed delay (which may be unreliable on a large network), we use central moments—which are not affected by the deterministic transmission delay.




Figure 1: Unidentifiability of Mean Delay: If one were to replace d_1 with $d_1 + \mu$ and d_2, d_3 with $d_2 - \mu, d_3 - \mu$ for $\mu > 0$ (assuming μ can be chosen so that all delays remain positive) then the distribution of delays at a, b would be unchanged. This example also shows that one cannot deduce the delays on all edges given total delays at all leaves.

1.2 Our Results.

Our main result is the following theorem.

Theorem 1 (Main Result) *Let $\varepsilon > 0$ and $J \geq 2$ be fixed. Let Assumptions 1 and 2 hold. Then, there is a polynomial-time algorithm which solves the multicast inference problem with high probability using $k = O(\text{poly}(\log n))$ samples.*

See Theorems 3, 4, and 5 below for more precise statements.

The proofs of the main theorems rely on the important notion of a tree metric from phylogenetics. Roughly speaking, a tree metric is a metric on the leaves of a tree such that the distance between any two leaves can be written as a sum of edge weights on the corresponding path. (See Section 2 for definitions.) There are two components to our algorithm:

1. **Topology reconstruction:** The reconstruction of the routing tree can be achieved by adapting known phylogenetic reconstruction algorithms—once the proper delay-based metric is defined. This result is proved in Section 3. The relevant phylogenetic background is introduced in Section 2.
2. **Moment estimation on edges:** Most of the technical work of this paper is in deriving and analyzing a metric-based algorithm for inferring edge delay distributions (Theorems 4 and 5). For this purpose, a) we relax the notion of a tree metric to allow nonnegative edge weights, b) we define appropriate delay-based metrics, and c) we show how to estimate these metrics. The analysis relies on large deviations arguments.

As far as we are aware, our algorithm is the first multicast inference algorithm to be both provably efficient and consistent. Previous work concerned mostly non-rigorous techniques such as maximum pseudo-likelihood and EM algorithms. See [6] for details. An exception is the independent, unpublished work of Liang et al. [15] which uses techniques similar to ours in the related context of multicast packet drop inference.

1.3 Discussion

Validity of assumptions. The multicast delay process defined in Section 1.1 relies on two basic assumptions about routing and traffic which makes its analysis possible: *temporal* and *spatial independence*. In reality, of course, both assumptions are violated to some extent. Lo Presti et al. [23] (see also [4]) studied the effect of these violations empirically and concluded that the multicast delay

process is a useful first approximation to the underlying complex process. We briefly summarize their findings.

Temporal dependence—delays at a given link being correlated at different points in time—is common in communication networks. But, as it turns out, its impact is rather mild for our purposes. Indeed the type of inference procedure studied in [23] (as well as in the current paper) does not actually require independence in time but only *ergodicity*—a much weaker assumption; more precisely, the estimator in [23] (and in the current paper) is consistent as long as the delay process is ergodic. Hence, the temporal dependencies impact only the *convergence rate* of the inference procedure. Lo Presti et al. showed empirically that, although this effect cannot be ignored, it is rather mild. Quantifying exactly the effect of temporal correlations on the theoretical convergence rate of an estimator is non-trivial.

As for *spatial correlations*—dependencies in delays on neighboring links—Lo Presti et al. found that they can produce a systematic bias in the estimation. However, they showed empirically that the bias is a small, second-order effect, possibly—they argue—because the diversity of traffic on the network results only in localized, short-term correlations in delays. They also point out that very little is known about the precise structure of such spatial correlations in real networks, making it hard to derive a good model for them.

Another assumption implicit in our model is that the process, including the routing tree itself, remains homogeneous over time. In fact, there are sporadic large-scale changes in the network. These explain why a *low sample complexity* is critical for an inference procedure to be useful in practice. Minimizing the sample complexity is the main focus of this paper.

Related results. The *multicast delay inference* problem was formalized by Lo Presti et al. in [23]. In that paper, the authors give a procedure to infer a discretized delay distribution on each link, given the routing tree topology. Their algorithm is based on an ad-hoc fixed point equation that is solved by least squares. Moreover, these authors show that their estimator is asymptotically normal with a variance-covariance matrix depending implicitly on the delay characteristics. More explicit formulas are given in the limit of small delays. The algorithm is tested on small networks and the dependence on the size is not given.

More recently, Ni and Tatikonda [19, 20, 21, 22]—in work subsequent to ours [2]—used phylogenetic techniques to recover the routing tree topology in this context. Similarly to the current paper, they use distance-based techniques. The basic algorithm they consider is the well-known Neighbor-Joining (NJ) algorithm which they apply to various tree metrics, for instance, the delay variance metric (as we do here). They also deal with trees of internal degrees higher than 3 by introducing a variant of NJ called Rooted Neighbor-Joining (RNJ) [21] (based on a technique equivalent to what is known in phylogenetics as the Farris transform [10]). They show more precisely that RNJ is a consistent estimator of the routing tree, but no convergence rate is given. Note, however, that RNJ has in fact a high sample complexity due to its reliance on the *diameter* of the tree. See, e.g., [1]. See also our discussion about diameter v. depth in Section 2.2. Here, we make use of state-of-art phylogenetic reconstruction techniques to derive a low sample complexity algorithm for routing tree reconstruction. We also show how to infer delay distributions. A technique to infer discrete delays was also subsequently obtained by Ni and Tatikonda [20] (although no convergence rate is provided).

A related network tomography problem is the so-called *multicast link loss inference* problem, where one observes packet losses at the receivers of a multicast routing tree—instead of delays—and seeks to infer the routing tree and packet drop probabilities on the links. This problem was formalized in [4] where a maximum-likelihood estimation procedure was analyzed. In [4], the network topology is assumed known. In more recent independent work, Liang et al. [15] (unpublished) applied phylogenetic techniques to the inference of the routing topology in this context. Indeed, the multicast link loss problem is in essence a special case of the standard model of DNA evolution used in biology. Similarly to the current paper, Liang et al. use distance-based techniques. More precisely, they give a com-

putationally efficient reconstruction algorithm with sample complexity $O(b^{-2} \log n)$ where b (possibly depending on n) is a lower bound on the link loss probability. Ni and Tatikonda [19, 20, 21, 22] (see above) also considered the link loss inference problem.

1.4 Organization of the Paper

The paper is organized as follows. We start with some phylogenetic background in Section 2. Our results concerning the topology reconstruction can be found in the Section 3. We then present and analyze our delay inference algorithm in Section 4.

2 Phylogenetic Reconstruction Techniques

In this section, we summarize and adapt to our setting the DMR algorithm of [8].

2.1 Basics

We begin with a few basic notions from phylogenetics.

Tree metrics. In phylogenetics, the notion of a tree metric is useful for reconstructing the topology of phylogenies. We use the notation $\mathbb{R}_{++} = \{x \in \mathbb{R} : x > 0\}$.

Definition 3 (Tree Metric) *Let L be a finite set with cardinality n . A function $W : L \times L \rightarrow \mathbb{R}_+$ defines a (nondegenerate) tree metric if the following holds. There exist a tree $T = (V, E)$ with leaf set L and a weight function $w : E \rightarrow \mathbb{R}_{++}$ such that $W(a, b) = \sum_{e \in P_{ab}} w_e$ for all $a, b \in L$ where P_{ab} is the path between a and b in T .*

Tree metrics are usually estimated from samples of the tree process at the leaves. In that context, Azuma’s inequality is useful (see, e.g., [18]).

Lemma 1 (Azuma-Hoeffding Inequality) *Suppose $X = (X_1, \dots, X_k)$ are independent random variables taking values in a set S , and $f : S^k \rightarrow \mathbb{R}$ is any t -Lipschitz function: $|f(\mathbf{x}) - f(\mathbf{y})| \leq t$ whenever \mathbf{x} and \mathbf{y} differ at just one coordinate. Then, $\forall \lambda > 0$,*

$$\mathbb{P}[f(X) - \mathbb{E}[f(X)] \geq \lambda] \leq \exp\left(-\frac{\lambda^2}{2t^2k}\right),$$

and

$$\mathbb{P}[f(X) - \mathbb{E}[f(X)] \leq -\lambda] \leq \exp\left(-\frac{\lambda^2}{2t^2k}\right).$$

Bipartitions. A useful combinatorial description of a tree $T = (V, E)$ is obtained by noticing that each edge $e \in E$ of the tree naturally corresponds to a partition of the leaves L into two subsets (that is, the leaves on either “side” of e). Such partitions are called *bipartitions* and they characterize the tree: it is easy to generate all bipartitions corresponding to a given tree, and on the other hand, there is a simple efficient iterative procedure to recover a tree from the set of all of its bipartitions. See [11, 25] for details.

2.2 Distorted Metric Algorithms

Classical distance-based reconstruction algorithms (that is, those methods based on tree metrics) such as UPGMA [26] or Neighbor-Joining (NJ) [24], typically make use of *all* pairwise distances between leaves. This leads to difficulties because “long” distances are more “noisy” and require a large number of samples to be accurately estimated. For instance, in the phylogenetic context, the widely used NJ algorithm is computationally efficient, but it is known to require *exponentially* many samples—even for simple linear trees [13].

An important breakthrough was made in [9] where it was shown that it was in fact enough to use “short” distances to fully recover the tree under reasonable assumptions. To help understand this result, we need a notion of tree “depth.” Given an edge $e \in E$, the *chord depth* of e is the length (in graph distance) of the shortest path between two leaves on which e lies¹. That is,

$$\Delta(e) = \min \{d(u, v) : u, v \in L, e \in P_{uv}\},$$

where d is the graph distance on T . We define the *chord depth of a tree T* to be the maximum chord depth in T

$$\Delta(T) = \max \{\Delta(e) : e \in E\}.$$

It is easy to show that $\Delta(T) \leq \log_2 n$ if the degree of all internal nodes is at least 3 (argue by contradiction). In a nutshell, the key insight behind the results in [9] is that the diameter and the depth of a tree behave very differently: even though the diameter can be as large as $O(n)$, the depth is always $O(\log n)$, in other words, each edge lies on a “short” path between two leaves. Using clever combinatorial arguments, Erdős et al. [9] showed that one can reconstruct trees with much fewer samples by ignoring those distances corresponding to paths longer than $O(\log n)$.

More recently, Daskalakis et al. [8] relaxed some of the assumptions in [9]. In particular, they gave a reconstruction algorithm based on short distances allowing internal degrees bigger than 3—which is particularly relevant in the networking context. Their algorithm, which we will call the DMR algorithm, reconstructs all bipartitions using only distances smaller than a threshold of order $O(\log n)$. To check that the algorithm works, one only needs to show that such distances are accurately estimated for a given number of samples. In the tomography setting, the DMR algorithm will allow us to reconstruct the routing tree using as few as $\text{poly log } n$ samples (see next section). The details of the algorithm are sketched in Appendix B.

We now state a corollary of [8] that will be useful to us. We first need the following definition which formalizes the idea that short distances are accurately estimated (and that long distances can in some sense be ignored).

Definition 4 (Distorted Metric [16, 12]) *Let $T = (V, E)$ be a tree with leaf set L and edge weight function $w : E \rightarrow \mathbb{R}_{++}$. Let $W : L \times L \rightarrow \mathbb{R}_+$ be the corresponding tree metric. Fix $\tilde{\tau}, \tilde{M} > 0$. We say that $\widehat{W} : L \times L \rightarrow (0, +\infty]$ is a $(\tilde{\tau}, \tilde{M})$ -distorted metric for T or a $(\tilde{\tau}, \tilde{M})$ -distortion of W if:*

1. (Symmetry) For all $u, v \in L$, \widehat{W} is symmetric, that is,

$$\widehat{W}(u, v) = \widehat{W}(v, u);$$

2. (Distortion) \widehat{W} is accurate on “short” distances, that is, for all $u, v \in L$, if either $W(u, v) < \tilde{M} + \tilde{\tau}$ or $\widehat{W}(u, v) < \tilde{M} + \tilde{\tau}$ then

$$|W(u, v) - \widehat{W}(u, v)| < \tilde{\tau}.$$

¹Note that unlike [8] we use the graph distance in the definition of chord depth. Because of our assumptions (see below) the two graph and weighted distances are the same up to a constant factor. Note also that we are using a different definition than [9]. But again the difference is only a constant factor.

Let $f, g > 0$ be bounds on the edge weights, that is, $f \leq w_e \leq g$ for all $e \in E$. We say that such an edge weight function satisfies the (f, g) -condition.

Theorem 2 (DMR Algorithm [8]) *Let $0 < f < g < +\infty$, $\tilde{\alpha} < 1/6$, and $\tilde{\beta} > 2$. There is a polynomial-time algorithm \mathbb{A} such that, for all trees $T = (V, E)$ with edge weight function w satisfying the (f, g) -condition and all $(\tilde{\alpha}f, \tilde{\beta}g\Delta(T))$ -distortions \widehat{W} of W (where W is the tree metric corresponding to w), \mathbb{A} applied to \widehat{W} returns T .*

Note that the previous theorem is a deterministic statement about distorted metrics. We show how to estimate such a distorted metric from random samples with high probability in Section 3.2.

3 Routing Tree Reconstruction

The goal of this section is to reconstruct efficiently the topology of the routing tree using Theorem 2.

3.1 Variance Metric

From Definition 3, one can define a tree metric by first choosing a tree—in our case, the routing tree—and then defining a weight function on its edges. Any positive quantity can serve as a weight. The important point is that one must be able to estimate the resulting tree metric from samples at the leaves. This governs the choice of the weight function.

Let $T = (V, E)$ be the (unknown) routing tree with leaf set L and consider the choice of weights

$$w_e^{(2)} = \text{Var}[d_e],$$

for all $e \in E$ and the corresponding tree metric

$$W^{(2)}(a, b) \equiv \sum_{e \in P_{ab}} \text{Var}[d_e],$$

for all $a, b \in L$. Our first task is to check that this metric can be estimated from samples at the leaves. Let a, b be leaves and consider the quantity $\delta_{ab}^{(2)} \equiv \text{Var}[D_a - D_b]$ (where recall from (1) that D_u is the delay at u). The delays D_a and D_b are observed at the leaves a and b respectively and therefore the variance of $D_a - D_b$ can be easily estimated. Moreover, we claim that the equality $\delta_{ab}^{(2)} = W^{(2)}(a, b)$ holds. Indeed, denote γ_{ab} the common ancestor of a and b , that is, the node at which all three paths P_{ab} , P_{0a} , and P_{0b} intersect (where we assume $a, b \neq 0$). Then, by independence of the edge delays, we have

$$\delta_{ab}^{(2)} = \text{Var}[D_a - D_b] = \text{Var} \left[\sum_{e \in P_{a\gamma_{ab}}} d_e - \sum_{e \in P_{\gamma_{ab}b}} d_e \right] = \sum_{e \in P_{a\gamma_{ab}}} \text{Var}[d_e] + \sum_{e \in P_{\gamma_{ab}b}} \text{Var}[d_e] = W^{(2)}(a, b).$$

Therefore, we can estimate $W^{(2)}$ by estimating $\delta^{(2)}$ at the leaves.

To estimate $\delta_{ab}^{(2)}$ from k samples, we use the standard unbiased estimator for the variance of $D_a - D_b$

$$\hat{\delta}_{ab}^{(2)} = \frac{1}{k-1} \sum_{i=1}^k \left[(D_a^i - D_b^i) - \hat{\delta}_{ab}^{(1)} \right]^2,$$

where

$$\hat{\delta}_{ab}^{(1)} = \frac{1}{k} \sum_{i=1}^k (D_a^i - D_b^i).$$

Below, we will need to show that $\hat{\delta}_{ab}^{(2)}$ is well concentrated around $\delta_{ab}^{(2)}$, which follows from the Azuma-Hoeffding inequality (see Lemma 1). The next lemmas provide the necessary Lipschitz condition.

Lemma 2 Suppose $X = (X_1, \dots, X_k)$ are independent random variables taking values in $[-B, B]$ with $k \geq 2$. Then, the variance estimator

$$s_X^2 = \frac{1}{k-1} \sum_{i=1}^k (X_i - \bar{X})^2 = \frac{1}{k(k-1)} \sum_{i < j} (X_i - X_j)^2,$$

where \bar{X} is the sample average, is $\frac{4B^2}{k}$ -Lipschitz.

Proof: Let X be as above and let Y differ from X in one coordinate. Then

$$|s_X^2 - s_Y^2| \leq \frac{1}{k(k-1)} \sum_{i < j} |(X_i - X_j)^2 - (Y_i - Y_j)^2| \leq \frac{1}{k} 4B^2.$$

■

We then get immediately the following.

Lemma 3 (Lipschitz Constant for Delay-based Metric) Say $\hat{\delta}_{ab}^{(2)}$ is computed with k samples. Then, $\hat{\delta}_{ab}^{(2)}$ is then $\frac{4|P_{ab}|^2 M^2}{k}$ -Lipschitz.

■

3.2 Inferring the routing tree

Equipped with a legitimate tree metric, we use the DMR algorithm to infer the topology. Here, we use Theorem 2 to prove that the routing tree can be inferred with $\text{poly log } n$ samples at the leaves. This is our main result for this section. The main technical difficulty (unlike the phylogenetic case) is in controlling the deviation of “long distances.” (See second part of the proof.) Fix $\tilde{\alpha}, \tilde{\beta}, f, g$ as in Theorem 2. Note that by assumption we have

$$f \leq w_e^{(2)} \leq g, \tag{2}$$

for all e with

$$g \equiv M^2.$$

Theorem 3 (Efficient Network Inference) Let $T = (V, E)$ be the (unknown) routing tree where edge delays satisfy Assumptions 1 and 2. Consider the tree metric $W^{(2)} = \delta^{(2)}$ and assume that the estimate $\widehat{W}^{(2)} = \hat{\delta}^{(2)}$ is computed using k samples at the leaves. Then, DMR returns the correct topology for T with probability $1 - o(1)$ if $k = \Omega(\log^5 n)$ (where the constant factor depends only on f, g), as n tends to $+\infty$.

Proof: Assume k is as stated above. We apply Theorem 2 and therefore only need to show that $\widehat{W}^{(2)}$ is a $(\tilde{\alpha}f, \tilde{\beta}g\Delta(T))$ -distortion of $W^{(2)}$ when $k = \Omega(\log^5 n)$.

Part 1. First, we must show that distances smaller than $\tilde{\beta}g\Delta(T) + \tilde{\alpha}f$ under $W^{(2)}$ are approximated within $\tilde{\alpha}f$. For reasons that will become clear below, we show instead that distances smaller than *twice* that amount are well approximated. Let a, b be a pair of leaves at distance at most $2\tilde{\beta}g\Delta(T) + 2\tilde{\alpha}f$. Let \mathcal{A} be the probability that, for all such pairs, $W^{(2)}$ is approximated within $\tilde{\alpha}f$. By our assumption (2), the number of edges on the path between a and b is at most

$$|P_{ab}| \leq (2\tilde{\beta}\Delta(T) + 1)\frac{g}{f},$$

where we used that $f < g$ and $\tilde{\alpha} < 1/6$. By Lemmas 1 and 3, we have

$$\mathbb{P} \left[\left| \delta_{ab}^{(2)} - \hat{\delta}_{ab}^{(2)} \right| \geq \tilde{\alpha} f \right] \leq 2 \exp \left(- \frac{(\tilde{\alpha} f)^2 k}{2[4(2\tilde{\beta}g\Delta(T) + 1)^2 \frac{g^2}{f^2} M^2]^2} \right) \leq \frac{1}{\text{poly}(n)},$$

from $\Delta(T) = O(\log n)$, $k = \Omega(\log^5 n)$, and the fact that f, g, M are constants. The notation $\text{poly}(n)$ means $O(n^K)$ for a K as large as we need as long as the constant factor in k is large enough. Since there are at most n^2 such pairs of leaves, we get $\mathcal{A} \leq \frac{1}{\text{poly}(n)}$.

Part 2. Let a, b be a pair of leaves at distance at least $2\tilde{\beta}g\Delta(T) + 2\tilde{\alpha}f$ under $W^{(2)}$. We now show that, for all such pairs, $\widehat{W}^{(2)}$ is at least $\tilde{\beta}g\Delta(T) + \tilde{\alpha}f$. Let \mathcal{B} be the probability of that event. Note first that from Azuma-Hoeffding (Lemma 1), it follows that for any pair of leaves a, b ,

$$\mathbb{P} \left[|(D_a - D_b) - \mathbb{E}[D_a - D_b]| \leq \sqrt{|P_{ab}|} \Theta(\sqrt{\log n}) \right] \geq 1 - \frac{1}{\text{poly}(n)}. \quad (3)$$

Let \mathcal{E} be the event that the inequality in square brackets in (3) holds for all k samples used to compute $\hat{\delta}_{ab}^{(2)}$. Then from Lemma 2, on \mathcal{E} , the Lipschitz constant of $\hat{\delta}_{ab}^{(2)}$ (as a function of the centered samples $(D_a^i - D_b^i) - \mathbb{E}[D_a - D_b]$) is $t = \frac{|P_{ab}|}{k} \Theta(\log n)$ and therefore, by Lemma 1 again,

$$\begin{aligned} \mathbb{P} \left[\hat{\delta}_{ab}^{(2)} \leq \tilde{\beta}g\Delta(T) + \tilde{\alpha}f \mid \mathcal{E} \right] &\leq \mathbb{P} \left[\hat{\delta}_{ab}^{(2)} \leq \frac{\mathbb{E}[\hat{\delta}_{ab}^{(2)} | \mathcal{E}]}{2} \mid \mathcal{E} \right] \\ &\leq \mathbb{P} \left[\mathbb{E}[\hat{\delta}_{ab}^{(2)} | \mathcal{E}] - \hat{\delta}_{ab}^{(2)} \geq \frac{\mathbb{E}[\hat{\delta}_{ab}^{(2)} | \mathcal{E}]}{2} \mid \mathcal{E} \right] \\ &\leq \exp \left(- \frac{(\mathbb{E}[\hat{\delta}_{ab}^{(2)} | \mathcal{E}]/2)^2}{2t^2 k} \right) \\ &\leq \exp \left(- \frac{k}{O(\log^2 n)} \right) \\ &\leq \frac{1}{\text{poly}(n)}, \end{aligned}$$

where we used $\delta_{ab}^{(2)} = \Theta(|P_{ab}|)$ and

$$(1 - o(1))\delta_{ab}^{(2)} \leq \mathbb{E}[\hat{\delta}_{ab}^{(2)} | \mathcal{E}] \leq (1 + o(1))\delta_{ab}^{(2)},$$

which follows from $\mathbb{E}[\hat{\delta}_{ab}^{(2)} | \mathcal{E}] \mathbb{P}[\mathcal{E}] + \mathbb{E}[\hat{\delta}_{ab}^{(2)} | \mathcal{E}^c] \mathbb{P}[\mathcal{E}^c] = \mathbb{E}[\delta_{ab}^{(2)}]$, $\mathbb{E}[\delta_{ab}^{(2)}] = \delta_{ab}^{(2)}$, $\mathbb{P}[\mathcal{E}^c] \leq 1/\text{poly}(n)$, and $\hat{\delta}_{ab}^{(2)} = O(n^2)$. Therefore, we have $\mathcal{B} \leq 1/\text{poly}(n)$.

Combining the two parts of the argument, we have shown that, except with $o(1)$ probability, $\widehat{W}^{(2)}$ is a $(\tilde{\alpha}f, \tilde{\beta}g\Delta(T))$ -distortion of $W^{(2)}$. Indeed, by Part 2 the pairs of leaves for which $\widehat{W}^{(2)} < \tilde{\beta}g\Delta(T) + \tilde{\alpha}f$ must have $W^{(2)} < 2\tilde{\beta}g\Delta(T) + 2\tilde{\alpha}f$ and such pairs satisfy the approximation guarantee required by the definition of a distorted metric by Part 1. Moreover, Part 1 implies in particular that pairs of leaves such that $W^{(2)} < \tilde{\beta}g\Delta(T) + \tilde{\alpha}f$ also satisfy the approximation guarantee. ■

4 Edge Delay Inference

In this section, we show how to estimate the characteristic moments of edge delays. In Section 3, we showed how to reconstruct the topology efficiently with high probability (see Theorem 3). Therefore, along with Assumptions 1 and 2, we make the following assumption.

Assumption 3 (Correct Reconstruction of Routing Tree) *We assume that the routing tree was correctly estimated. (This is true with high probability by Theorem 3.)*

Our general idea to recover delay distributions is to define so-called “additive functions” whose edge weights are moments of delays. Then we use the AFI algorithm below to recover the moments efficiently from the data at the leaves. As it turns out, even moments are rather straightforward to estimate inductively while odd moments are trickier. Also, as in the tree reconstruction algorithm (see also [9, 17]), the AFI algorithm uses only “short” paths during the estimation process, which allows a significant reduction in the sample size (see Propositions 1, 2 and Theorems 4, 5 for details).

4.1 Additive Functions

In the remainder of this paper, we use additive metric-type ideas to estimate moments of edge delays. For this purpose, we need to recover edge weights from appropriately defined tree metrics. In fact, we use a notion of “generalized” tree metric which is useful in treating odd moments. This definition allows for negative edge weights.

Definition 5 (Additive function) *A function on the leaf set of the tree $W : L \times L \rightarrow \mathbb{R}$ is called an additive function on the leaves if there exists weights $w_e \in \mathbb{R}$ on each of the edges (not necessarily positive), such that for all leaves a, b*

$$W(a, b) = \sum_{e \in P_{ab}} w_e.$$

Suppose we are given access to an additive function W on the leaves. Our goal is now to recover the w_e ’s from the function W , assuming further that we are given the tree T . For this purpose, we use a standard algorithm from combinatorial phylogenetics—related to the so-called Four-Point Method of Buneman [3] (see also [11, 25]). We will refer to this algorithm as the ADDITIVE FUNCTION INFERENCE (AFI) algorithm. See Figures 2 and 3.

Algorithm ADDITIVE FUNCTION INFERENCE

Input: tree T , function W at the leaves;

Output: edge weights w_e , for all $e \in E$;

- For all internal edges e ,
 - Let S_1, \dots, S_4 be the four subtrees hanging from e as in Figure 3;
 - For each S_i , compute u_i the closest (in graph distance) leaf to the root r_i of S_i ;
 - Compute

$$w_e = \frac{1}{2}(W(u_1, u_3) + W(u_2, u_4) - W(u_1, u_2) - W(u_3, u_4)).$$

- For all leaf edges e ,
 - Let $e = (a, v)$ with a a leaf;
 - Proceed as above where u_3 and u_4 are set to a .

Figure 2: Algorithm ADDITIVE FUNCTION INFERENCE.

4.2 Delay-based metrics

Let $T = (V, E)$ be the routing tree with leaf set L and consider again the choice of weights

$$w_e^{(2)} = \text{Var}[d_e],$$




Figure 3: Edge weight inference.

for all $e \in E$ and

$$W^{(2)}(a, b) = \sum_{e \in P_{ab}} w_e^{(2)},$$

for all $a, b \in L$. Recall that

$$\delta_{ab}^{(2)} = \text{Var}[D_a - D_b] = \text{Var} \left[\sum_{e \in P_{a\gamma_{ab}}} d_e - \sum_{e \in P_{\gamma_{ab}b}} d_e \right] = \sum_{e \in P_{a\gamma_{ab}}} \text{Var}[d_e] + \sum_{e \in P_{\gamma_{ab}b}} \text{Var}[d_e] = W^{(2)}(a, b).$$

Therefore, using the AFI algorithm, we can recover estimates of the $w_e^{(2)}$'s from the $\hat{\delta}_{ab}^{(2)}$'s.

More generally, we let

$$w_e^{(j)} = \mathbb{E} \left[(\bar{d}_e)^j \right],$$

for all $e \in E$ where

$$\bar{d}_e = d_e - \mathbb{E}[d_e].$$

Also, let

$$W^{(j)}(a, b) = \sum_{e \in P_{ab}} w_e^{(j)},$$

for all $a, b \in L$. Let

$$\bar{D}_a = D_a - \mathbb{E}[D_a],$$

for all $a \in L$. Again, to obtain $W^{(j)}(a, b)$, we seek to use the quantity

$$\delta_{ab}^{(j)} = \mathbb{E} \left[(\bar{D}_a - \bar{D}_b)^j \right],$$

for $j > 1$, which can be estimated from the samples using

$$\hat{\delta}_{ab}^{(j)} = \frac{1}{k} \sum_{i=1}^k \left((D_a^i - D_b^i) - \hat{\delta}_{ab}^{(1)} \right)^j$$

where

$$\hat{\delta}_{ab}^{(1)} = \frac{1}{k} \sum_{i=1}^k (D_a^i - D_b^i).$$

As Lemma 4 below shows, this can be done inductively. However, the lemma also shows that odd moments have to be treated more carefully.

4.3 Algorithm for Moment Inference

We first need the following definitions. Let a, b be leaves and $j \in \mathbb{N}$. We use the notation $[h] = \{0, \dots, h\}$ for $h \in \mathbb{N}$. Recall that γ_{ab} is the most recent common ancestor of a and b in the tree. Denote $\nu = |P_{ab}|$, $\alpha = |P_{a\gamma_{ab}}|$, and $\beta = |P_{\gamma_{ab}b}|$, and define

$$\mathcal{D}_j(a, b) = \left\{ (\mathbf{x}, \mathbf{y}) \in [j-1]^\alpha \times [j-1]^\beta : \sum_{i=1}^\alpha x_i + \sum_{i=1}^\beta y_i = j \right\}.$$

For $(\mathbf{x}, \mathbf{y}) \in \mathcal{D}_j(a, b)$, let

$$\binom{j}{\mathbf{x}, \mathbf{y}} = \frac{j!}{\prod_{i=1}^\alpha x_i! \prod_{i=1}^\beta y_i!}.$$

and consider the function

$$\mathcal{F}_j(a, b) = \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}_j(a, b)} \binom{j}{\mathbf{x}, \mathbf{y}} \prod_{i=1}^\alpha w_{e_i}^{(x_i)} \prod_{i=1}^\beta (-1)^{y_i} w_{f_i}^{(y_i)},$$

where $P_{a\gamma_{ab}} = (e_1, \dots, e_\alpha)$ and $P_{\gamma_{ab}b} = (f_1, \dots, f_\beta)$.

Lemma 4 *Let $j \in \mathbb{N}$ and define the function $\mathcal{F}_j : L \times L \rightarrow \mathbb{R}$ as above. Then,*

1. *we have for all $a, b \in L$*

$$\delta_{ab}^{(j)} - \mathcal{F}_j(a, b) = \sum_{i=1}^\alpha w_{e_i}^{(j)} + (-1)^j \sum_{i=1}^\beta w_{f_i}^{(j)}, \quad (4)$$

2. *in particular, if j is even, we have for all $a, b \in L$*

$$\delta_{ab}^{(j)} - \mathcal{F}_j(a, b) = W^{(j)}(a, b). \quad (5)$$

Proof: This follows immediately from a multinomial expansion. ■

The important point to note in (5) is that $\mathcal{F}_j(a, b)$ depends only on delay moments *of order strictly less than j* and that $\delta_{ab}^{(j)}$ can be estimated from samples at the leaves. Therefore, if j is even and if we have estimates of all edge delay moments of order up to $j-1$, we can estimate $W^{(j)}(a, b)$ by (5). Using the AFI algorithm, we can then get an estimate of the j -th moments $w_e^{(j)}$. However, if j is odd, the coefficient $(-1)^j$ in (4) precludes the use of this procedure. Lemma 5 below shows how to handle this case. We note in passing that Lemma 4 above is sufficient for delay distributions symmetric about their mean. Indeed, in that case, all odd central moments are zero and one can use (5) recursively to estimate all even characteristic moments. See Figure 4.

We now tackle odd moments. A proper estimation procedure follows from the next lemma. We first need a few definitions. For $a, b \in L$, and $1 \leq i^* \leq \alpha$, we let

$$\mathcal{E}_j^{(1)}(a, b; i^*) = \left\{ (\mathbf{x}, \mathbf{y}) \in [j-1]^\alpha \times [j-1]^\beta : \sum_{i=1}^\alpha x_i + \sum_{i=1}^\beta y_i = j, x_{i^*} \geq 1 \right\}.$$

Algorithm SYMMETRIC EDGE RECONSTRUCTION

Input: data $\{D_a^1\}_{a \in L}, \dots, \{D_a^k\}_{a \in L}$ at the leaves; topology T ;

Output: estimated characteristic (even) moments $\hat{w}_e^{(j)}$ for all $e \in E$ and $2 \leq j \leq J$ even;

- Initialization: set all estimates of odd moments to 0;
- Main Loop: For all $2 \leq j \leq J$ even,

- For all $a, b \in L$,
 - * Estimate $\hat{\delta}_{ab}^{(j)}$;
 - * Estimate $\mathcal{F}_j(a, b)$ with

$$\hat{\mathcal{F}}_j(a, b) = \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}_j(a, b)} \binom{j}{\mathbf{x}, \mathbf{y}} \prod_{i=1}^{\alpha} \hat{w}_{e_i}^{(x_i)} \prod_{i=1}^{\beta} (-1)^{y_i} \hat{w}_{f_i}^{(y_i)}.$$

- * Compute

$$\widehat{W}^{(j)}(a, b) = \hat{\delta}_{ab}^{(j)} - \hat{\mathcal{F}}_j(a, b)$$

- Use the AFI algorithm on $\widehat{W}^{(j)}(a, b)$ to recover all $\hat{w}_e^{(j)}$'s.

Figure 4: Algorithm SYMMETRIC EDGE RECONSTRUCTION.

and

$$\mathcal{G}_j^{(1)}(a, b) = \sum_{i^*=1}^{\alpha} \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{E}_j^{(1)}(a, b; i^*)} x_{i^*} \binom{j-1}{\mathbf{x}, \mathbf{y}} \prod_{i=1}^{\alpha} w_{e_i}^{(x_i)} \prod_{i=1}^{\beta} (-1)^{y_i} w_{f_i}^{(y_i)},$$

where we use the notations of Lemma 4. Similarly, for $1 \leq i^* \leq \beta$, we define $\mathcal{E}_j^{(2)}(a, b; i^*)$ and $\mathcal{G}_j^{(2)}(a, b)$ by interchanging the roles of \mathbf{x} and \mathbf{y} . Our next definition requires a few combinatorial notions. Recall the definition of quartet split from Section 2. Let a, b, c be any leaves in a *rooted* tree T with root 0 (which is also a leaf). We write $ab|c$ if $ab|c0$ holds in T . Then, for all leaves $a, b, c \neq 0$ with $ab|c$, let

$$\phi_{ab|c}^{(j)} = \mathbb{E} \left[(\overline{D}_a - \overline{D}_b)^{j-1} (\overline{D}_a + \overline{D}_b - 2\overline{D}_c) \right].$$

Lemma 5 *Let $j \in \mathbb{N}$. Then, using the notations above, we have for all $a, b, c \in L$*

$$W^{(j)}(a, b) = \phi_{ab|c}^{(j)} - \left[\mathcal{G}_j^{(1)}(a, b) + \mathcal{G}_j^{(2)}(a, b) \right]. \quad (6)$$

Proof: We write

$$\mathbb{E} \left[(\overline{D}_a - \overline{D}_b)^{j-1} (\overline{D}_a + \overline{D}_b - 2\overline{D}_c) \right] = \mathbb{E} \left[(\overline{D}_a - \overline{D}_b)^{j-1} (\overline{D}_a - \overline{D}_c) \right] + \mathbb{E} \left[(\overline{D}_a - \overline{D}_b)^{j-1} (\overline{D}_b - \overline{D}_c) \right]$$

Let (as in Figure 5)

$$H_1 = \sum_{e \in P_{a\gamma_{ab}}} \bar{d}_e \quad H_2 = \sum_{e \in P_{b\gamma_{ab}}} \bar{d}_e \quad H_3 = \sum_{e \in P_{a\gamma_{ac}}} \bar{d}_e \quad H_4 = \sum_{e \in P_{c\gamma_{ac}}} \bar{d}_e.$$




Figure 5: The H_i 's are centered sums of delays on the corresponding paths.

Note that all these random variables are independent and have 0 mean. Then

$$\begin{aligned}
\mathbb{E} \left[(\bar{D}_a - \bar{D}_b)^{j-1} (\bar{D}_a - \bar{D}_c) \right] &= \mathbb{E} \left[(H_1 - H_2)^{j-1} (H_1 + H_3 - H_4) \right] \\
&= \mathbb{E} \left[(H_1 - H_2)^{j-1} (H_1) \right] \\
&= \mathbb{E} \left[\left(\sum_{e \in P_{a\gamma_{ab}}} \bar{d}_e - \sum_{e \in P_{b\gamma_{ab}}} \bar{d}_e \right)^{j-1} \left(\sum_{e \in P_{a\gamma_{ab}}} \bar{d}_e \right) \right] \\
&= \sum_{e \in P_{a\gamma_{ab}}} w_e^{(j)} + \mathcal{G}_j^{(1)}(a, b).
\end{aligned}$$

Similarly,

$$\mathbb{E} \left[(\bar{D}_a - \bar{D}_b)^{j-1} (\bar{D}_b - \bar{D}_c) \right] = \sum_{e \in P_{b\gamma_{ab}}} w_e^{(j)} + \mathcal{G}_j^{(2)}(a, b).$$

The result follows. ■

Again, the key point in (6) is that $\mathcal{G}_j^{(1)}(a, b)$ and $\mathcal{G}_j^{(2)}(a, b)$ depend only on moments of order *strictly less than* j and that $\phi_{ab|c}^{(j)}$ can be estimated from samples at the leaves. The algorithm for the general case is detailed in Figure 6. We use the plugin estimator for $\phi_{ab|c}^{(j)}$,

$$\hat{\phi}_{ab|c}^{(j)} = \frac{1}{k} \sum_{i=1}^k \left((D_a^i - D_b^i) - \hat{\delta}_{ab}^{(1)} \right)^{j-1} \left((D_a^i - D_c^i) - \hat{\delta}_{ac}^{(1)} + (D_b^i - D_c^i) - \hat{\delta}_{bc}^{(1)} \right).$$

5 Analysis of the ER Algorithm

We start with the analysis of the symmetric case.

We begin with a concentration result for the estimate $\hat{\delta}_{ab}^{(j)}$. For convenience, we assume $M \geq 1$. (This can always be obtained by rescaling.) Recall the definition of the depth of T from Section 2 and remember that $\Delta(T) = O(\log n)$ if the degree of all internal nodes is at least 3. The dependence of our bounds on the depth of the routing tree explains the importance of using short paths in the estimation procedures.

Algorithm EDGE RECONSTRUCTION

Input: data $\{D_a^1\}_{a \in L}, \dots, \{D_a^k\}_{a \in L}$ at the leaves; topology T ;

Output: estimated characteristic moments $\hat{w}_e^{(j)}$ for all $e \in E$ and $2 \leq j \leq J$;

- Initialization: set all estimates of first moments to 0;
- Main Loop: For all $2 \leq j \leq J$,
 - For all $a, b \in L$,
 - * Pick the closest leaf c above γ_{ab}
 - * Compute $\hat{\phi}_{ab|c}^{(j)}$, the plug-in estimator for $\phi_{ab|c}^{(j)}$;
 - * Estimate $\mathcal{G}_j^{(1)}(a, b)$ with

$$\hat{\mathcal{G}}_j^{(1)}(a, b) = \sum_{i^*=1}^{\alpha} \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{E}_j^{(1)}(a, b; i^*)} x_{i^*} \binom{j-1}{\mathbf{x}, \mathbf{y}} \prod_{i=1}^{\alpha} \hat{w}_{e_i}^{(x_i)} \prod_{i=1}^{\beta} (-1)^{y_i} \hat{w}_{f_i}^{(y_i)},$$

and similarly for $\mathcal{G}_j^{(2)}(a, b)$;

- * Compute

$$\widehat{W}^{(j)}(a, b) = \hat{\phi}_{ab|c}^{(j)} - \left(\hat{\mathcal{G}}_j^{(1)}(a, b) + \hat{\mathcal{G}}_j^{(2)}(a, b) \right),$$

- Use the AFI algorithm on $\widehat{W}^{(j)}(a, b)$ to recover all $\hat{w}_e^{(j)}$'s.

Figure 6: Algorithm EDGE RECONSTRUCTION.

Proposition 1 *Let $a, b \in L$ at graph distance less than 2Δ where $\Delta = \Delta(T)$ is the chord depth of T . Fix $j \in \mathbb{N}$. We have the following (where the constants depend on J and M only):*

1. *There exists a constant C such that, $\forall \lambda > 0$,*

$$\mathbb{P} \left(\left| \hat{\delta}_{ab}^{(j)} - \mathbb{E} \left[\hat{\delta}_{ab}^{(j)} \right] \right| > \lambda \right) \leq 2 \exp \left(-\frac{\lambda^2 k}{C \Delta^{2j-1}} \right). \quad (7)$$

2. *There exists a constant C' such that*

$$\mathbb{E} \left(\left| \delta_{ab}^{(1)} - \hat{\delta}_{ab}^{(1)} \right|^j \right) \leq C' \frac{\Delta^j}{k^{j/2}}, \quad (8)$$

where $\delta_{ab}^{(1)} = \mathbb{E}[D_a - D_b]$.

3. *There exists a constant C'' such that, if $k \geq \Delta^2$,*

$$\left| \mathbb{E} \left[\hat{\delta}_{ab}^{(j)} \right] - \delta_{ab}^{(j)} \right| \leq C'' \frac{M^{2j} \Delta^{j+1}}{\sqrt{k}}. \quad (9)$$

4. *If further*

$$C'' \frac{M^{2j} \Delta^{j+1}}{\sqrt{k}} \leq \lambda,$$

then we have

$$\mathbb{P} \left[\left| \hat{\delta}_{ab}^{(j)} - \delta_{ab}^{(j)} \right| > 2\lambda \right] \leq 2 \exp \left(-\frac{\lambda^2 k}{C \Delta^{2j-1}} \right). \quad (10)$$

Proof: 1. We use Azuma's inequality (see Lemma 1). Let

$$\mathcal{K}_i = (D_a^i - D_b^i) - \hat{\delta}_{ab}^{(1)},$$

where D_u^i is the i -th delay sample at node u . Because $|P_{ab}| \leq 2\Delta$ and $d_e \in [0, M]$ for all e , it follows that

$$|\mathcal{K}_i| \leq 4M\Delta.$$

Then let

$$\mathcal{L} = \frac{1}{k} \sum_{i=1}^k (\mathcal{K}_i)^j,$$

and let \mathcal{L}' be the same quantity when an arbitrary d_e^i is perturbed by δ with $|\delta| \leq M$ (where d_e^i is the i -th delay sample on edge e). Without loss of generality, assume the perturbation is in the first sample. Then,

$$\mathcal{L}' = \frac{1}{k} \left(\left(\mathcal{K}_1 + \frac{(k-1)}{k} \delta \right)^j + \sum_{i=2}^k \left(\mathcal{K}_i - \frac{\delta}{k} \right)^j \right). \quad (11)$$

Now expanding (11), we get

$$\begin{aligned} |\mathcal{L} - \mathcal{L}'| &\leq \frac{1}{k} \left(2^j (4M\Delta)^{j-1} M + (k-1) \left(2^j (4M\Delta)^{j-1} \frac{M}{k} \right) \right) \\ &\leq C \frac{\Delta^{j-1}}{k}, \end{aligned}$$

for some constant C depending on M, J . Noting that \mathcal{L} depends on at most $2\Delta k$ random variables d_e^i , we get the result by an application of Azuma's inequality (for a different C).

2. Note that

$$\mathcal{L} = \delta_{ab}^{(1)} - \hat{\delta}_{ab}^{(1)},$$

is a $\frac{2M\Delta}{k}$ -Lipschitz function of $\{D_a^i - D_b^i\}_{i \in [k]}$ thus we have by Azuma's inequality

$$\mathbb{P} \left[\left| \delta_{ab}^{(1)} - \hat{\delta}_{ab}^{(1)} \right| > \lambda \right] \leq 2 \exp \left(-\frac{k\lambda^2}{8M^2\Delta^2} \right).$$

Now we use the fact that for a positive random variable Y ,

$$\mathbb{E} [Y^j] = j \int_0^\infty \lambda^{j-1} \mathbb{P}(Y > \lambda) d\lambda.$$

If $Y = |\delta_{ab}^{(1)} - \hat{\delta}_{ab}^{(1)}|$ and $\psi = \frac{k}{8M^2\Delta^2}$, we have

$$\mathbb{E} [Y^j] \leq \psi^{-\frac{j}{2}} \int_0^{+\infty} y^{\frac{j}{2}-1} e^{-y} dy = \left(\frac{8M^2\Delta^2}{k} \right)^{j/2} C'.$$

That proves 2 (for a different C').

3. We have

$$\hat{\delta}_{ab}^{(j)} = \frac{1}{k} \sum_{i=1}^k \left((D_a^i - D_b^i) - \hat{\delta}_{ab}^{(1)} \right)^j = \frac{1}{k} \sum_{i=1}^k \left((D_a^i - D_b^i - \delta_{ab}^{(1)}) + (\delta_{ab}^{(1)} - \hat{\delta}_{ab}^{(1)}) \right)^j.$$

Now expand using the binomial theorem and take expectations to get

$$\begin{aligned} \left| \mathbb{E} \left[\hat{\delta}_{ab}^{(j)} \right] - \delta_{ab}^{(j)} \right| &\leq \frac{1}{k} \mathbb{E} \left| \sum_{i=1}^k \sum_{h=0}^{j-1} \binom{j}{h} \left(D_a^i - D_b^i - \delta_{ab}^{(1)} \right)^h (\hat{\delta}_{ab}^{(1)} - \delta_{ab}^{(1)})^{j-h} \right| \\ &\leq C'' (4M\Delta)^j \max_{0 \leq h \leq j-1} \left\{ \mathbb{E} \left| \delta_{ab}^{(1)} - \hat{\delta}_{ab}^{(1)} \right|^{j-h} \right\}. \end{aligned}$$

Note that by $k \geq \Delta^2$, it follows that the maximum is attained at $h = j - 1$ in (8).

4. This follows from 1. and 3. ■

We then get the main theorem in the symmetric case. Recall that $J = O(1)$ and that, in general, $\Delta = O(\log n)$ where n is the number of leaves.

Theorem 4 *Let $\varepsilon > 0$ be arbitrarily small. If $k = \omega(\Delta^{2J^2} \log n)$, then after an application of SYMER, one has*

$$\mathbb{P} \left[\left| \hat{\delta}_e^{(j)} - \delta_e^{(j)} \right| \leq \varepsilon, \forall e \in E, \forall 1 \leq j \leq J \right] \geq 1 - o(1), \quad (12)$$

as $n \rightarrow +\infty$. The algorithm runs in time $O(\Delta^J n^2)$.

Proof: Let $(a, b) \in L \times L$ be called a *short pair* if a, b are at graph distance at most 2Δ . Denote \mathcal{S} be the set of all short pairs. Let

$$\sigma_j = \max_{(a,b) \in \mathcal{S}} \left| \widehat{W}^{(j)}(a, b) - W^{(j)}(a, b) \right|,$$

and

$$\Sigma_j = \max_{1 \leq i \leq j} \sigma_i.$$

It follows immediately from the application of the AFI algorithm that

$$\max_{e \in E} \left| \hat{w}_e^{(j)} - w_e^{(j)} \right| \leq 2\sigma_j.$$

Therefore, it suffices to prove

$$\Sigma_J = o(1),$$

with high probability as n tends to $+\infty$.

Further, assume we have a uniform bound

$$\max_{1 \leq j \leq J} \max_{(a,b) \in \mathcal{S}} \left| \hat{\delta}_{ab}^{(j)} - \delta_{ab}^{(j)} \right| \leq \tau^*.$$

Recall that

$$\widehat{W}^{(j)}(a, b) = \hat{\delta}_{ab}^{(j)} - \widehat{\mathcal{F}}_j(a, b)$$

where

$$\widehat{\mathcal{F}}_j(a, b) = \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}_j(a, b)} \binom{j}{\mathbf{x}, \mathbf{y}} \prod_{i=1}^{\alpha} \hat{w}_{e_i}^{(x_i)} \prod_{i=1}^{\beta} (-1)^{y_i} \hat{w}_{f_i}^{(y_i)}.$$

Note that $\widehat{\mathcal{F}}_j(a, b)$ has at most Δ^j terms (including the multinomial factor). Therefore, since the function

$$h(\mathbf{x}) = \prod_{j=1}^J x_j,$$

is continuously differentiable with bounded derivatives in $[-M^J, M^J]$, there is C (depending on M, J) such that

$$\sigma_j \leq \tau^* + C\Delta^j(2\Sigma_{j-1}),$$

for small Σ_{j-1} . Then we have

$$\Sigma_J \leq \tau^* C^* \Delta^{J^2/2},$$

for some $C^* > 0$ depending on J, M , where we used $\sigma_2 \leq \tau^*$.

So it suffices to have $\tau^* = (\omega_n \Delta^{J^2/2})^{-1}$ where $\omega_n \rightarrow +\infty$ as $n \rightarrow +\infty$ arbitrarily slowly. By the last part of Proposition 1, using a union bound over the $O(n^2)$ short pairs of leaves, it follows that $k = C' \omega_n \Delta^{2J^2} \log n$ samples are enough to guarantee

$$\mathbb{P} \left[\left| \hat{\delta}_{ab}^{(j)} - \delta_{ab}^{(j)} \right| \leq (\omega_n \Delta^{J^2/2})^{-1}, \forall 1 \leq j \leq J, \forall \text{ short pairs } a, b \right] \geq 1 - o(1),$$

for some C' depending on J, M .

As for the computational complexity of the algorithm, assume first that the tree is represented in such a way that finding the set of edges on the path between two leaves a, b at distance $O(\Delta)$ takes time $O(\Delta)$ (this is easy in a rooted tree). Note that for each j, a, b the sum

$$\hat{\mathcal{F}}_j(a, b) = \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}_j(a, b)} \binom{j}{\mathbf{x}, \mathbf{y}} \prod_{i=1}^{\alpha} \hat{w}_{e_i}^{(x_i)} \prod_{i=1}^{\beta} (-1)^{y_i} \hat{w}_{f_i}^{(y_i)}.$$

can be computed in time Δ^J . Since there are $O(n^2)$ pairs of leaves, the total complexity is $O(\Delta^J n^2)$. \blacksquare

Similarly, in the general case, we get:

Proposition 2 *Let $a, b, c \in L$ at graph distance less than 2Δ where $\Delta = \Delta(T)$ is the depth of T . Fix $j \in \mathbb{N}$. We have the following (where the constants depend on J and M only):*

1. *There exists a constant C such that, $\forall \lambda > 0$,*

$$\mathbb{P} \left(\left| \hat{\phi}_{ab|c}^{(j)} - \mathbb{E} \left[\hat{\phi}_{ab|c}^{(j)} \right] \right| > \lambda \right) \leq 2 \exp \left(-\frac{\lambda^2 k}{C \Delta^{2j-1}} \right). \quad (13)$$

2. *There exists a constant C' such that, if $k \geq \Delta^2$,*

$$\left| \mathbb{E} \left[\hat{\phi}_{ab|c}^{(j)} \right] - \phi_{ab|c}^{(j)} \right| \leq C' \frac{(M\Delta)^j}{\sqrt{k}}. \quad (14)$$

3. *If further*

$$C' \frac{(M\Delta)^j}{\sqrt{k}} \leq \lambda,$$

then we have

$$\mathbb{P} \left(\left| \hat{\phi}_{ab|c}^{(j)} - \phi_{ab|c}^{(j)} \right| > 2\lambda \right) \leq 2 \exp \left(-\frac{\lambda^2 k}{C \Delta^{2j-1}} \right). \quad (15)$$

Proof Sketch: The proof is very similar to Proposition 1. We only give a sketch.

To prove 1., it is enough to consider four separate cases depending on which path segment (corresponding to H_1, H_2, H_3 and H_4 in Figure 5) we make the perturbation.

To prove 2., note that we can write

$$\hat{\phi}_{ab|c} = \frac{1}{k} \sum_{\mathbf{I}}^k (X_{\mathbf{I}} + \epsilon_{\mathbf{I}})^{j-1} (Y_{\mathbf{I}} + \epsilon_{\mathbf{I}}), \quad (16)$$

with $X_i = (D_a^i - D_b^i) - \delta_{ab}^{(1)}$, $Y_i = (D_a^i - D_c^i - \delta_{ac}^{(1)}) + (D_b^i - D_c^i - \delta_{bc}^{(1)})$, and

$$\begin{aligned}\epsilon_I &= \delta_{ab}^{(1)} - \hat{\delta}_{ab}^{(1)}, \\ \epsilon_{II} &= (\delta_{ac}^{(1)} - \hat{\delta}_{ac}^{(1)}) + (\delta_{bc}^{(1)} - \hat{\delta}_{bc}^{(1)}).\end{aligned}$$

Also note that $\phi_{ab|c}^{(j)} = \mathbb{E}[X_i^{j-1}Y_i], \forall i$. Use the Binomial theorem to expand the expression in (16) and write it as

$$\hat{\phi}_{ab|c}^{(j)} = \frac{1}{k} \sum_{i=1}^k X_i^{j-1} Y_i + \mathcal{R},$$

where the error term is

$$\mathcal{R} = \epsilon_{II} \frac{1}{k} \sum_{i=1}^k (X_i + \epsilon_I)^{j-1} + \frac{1}{k} \sum_{i=1}^k Y_i \sum_{l=1}^{j-1} \binom{j-1}{l} \epsilon_I^l X_i^{j-1-l}.$$

Now use the fact that $|X_i| \leq 4M\Delta$, $|Y_i| \leq 8M\Delta$, and Part 2. of Proposition 1 to conclude that

$$\mathbb{E}[|\mathcal{R}|] \leq C' \frac{(M\Delta)^j}{\sqrt{k}}.$$

Part 3. now follows by combining Part 1. and 2. ■

Theorem 5 *Let $\varepsilon > 0$ be arbitrarily small. If $k = \omega(\Delta^{2J^2} \log n)$, then after an application of ER, one has*

$$\mathbb{P} \left[\left| \hat{\delta}_e^{(j)} - \delta_e^{(j)} \right| \leq \varepsilon, \forall e \in E, \forall 1 \leq j \leq J \right] \geq 1 - o(1), \quad (17)$$

as $n \rightarrow +\infty$. The algorithm runs in time $O(\Delta^J n^2)$.

Proof: The proof is identical to Theorem 4. ■

6 Concluding Remarks

1. We have assumed that delays are *finitely supported*. This assumption is not essential. Unbounded distributions for which similar concentration inequalities can be obtained lead to the same results. For example, using [14, Proposition 4.18], one can treat the case of Exponential and Gamma delays.
2. It is an interesting problem, from a practical point of view, to improve the dependence of our results on J .
3. It is somewhat intriguing that the reconstruction of the topology of the tree required the joint distributions on *pairs* of leaves whereas the reconstruction of delays (in the asymmetric case) required the joint distributions on *triples* of leaves. A similar situation holds in phylogenetics [7]. It could be interesting to prove that this is indeed necessary in some sense.
4. Throughout, the model was assumed to be static. In real-life networks, characteristics of the network change over time. One could try to adapt our algorithm to a more dynamic setting. See for example [5] for a discussion of temporal issues.

Acknowledgments

We thank Gang Liang, Elchanan Mossel, and Bin Yu for discussions and encouragements. S.R. gratefully acknowledges the partial support of CIPRES (NSF ITR grant # NSF EF 03-31494), NSERC (Canada), FQRNT (Quebec, Canada), a Loève Fellowship (UC Berkeley), and NSF grant DMS-0528488. S.R. also thanks Martin Nowak and the Program for Evolutionary Dynamics at Harvard University, where part of this work was done. S.B. thanks David Aldous for interesting discussions.

References

- [1] K. Atteson, *The performance of neighbor-joining methods of phylogenetic reconstruction*, *Algorithmica* **25** (1999), no. 2-3, 251–278. MR MR1703580 (2000k:92013)
- [2] S. Bhamidi, R. Rajagopal, and S. Roch, *Network delay inference from additive metrics*, Preprint. Available at Arxiv: math.PR/0604367, 2006.
- [3] P. Buneman, *The recovery of trees from measures of dissimilarity*, *Mathematics in the Archaeological and Historical Sciences*, Edinburgh University Press, Edinburgh, 1971, pp. 187–395.
- [4] Ramón Cáceres, N. G. Duffield, Joseph Horowitz, and Donald F. Towsley, *Multicast-based inference of network-internal loss characteristics*, *IEEE Trans. Inform. Theory* **45** (1999), no. 7, 2462–2480. MR MR1725131
- [5] Jin Cao, Drew Davis, Scott Vander Wiel, and Bin Yu, *Time-varying network tomography: router link data*, *J. Amer. Statist. Assoc.* **95** (2000), no. 452, 1063–1075. MR MR1821715
- [6] Rui Castro, Mark Coates, Gang Liang, Robert Nowak, and Bin Yu, *Network tomography: recent developments*, *Statist. Sci.* **19** (2004), no. 3, 499–517. MR MR2185628
- [7] Joseph T. Chang, *Full reconstruction of Markov models on evolutionary trees: identifiability and consistency*, *Math. Biosci.* **137** (1996), no. 1, 51–73. MR MR1410044 (97k:92011)
- [8] Constantinos Daskalakis, Elchanan Mossel, and Sébastien Roch, *Phylogenies without branch bounds: Contracting the short, pruning the deep*, To appear in RECOMB’09. Preprint available as arXiv:0801.4190v1, 2009.
- [9] Péter L. Erdős, Michael A. Steel, László A. Székely, and Tandy J. Warnow, *A few logs suffice to build (almost) all trees. I*, *Random Structures Algorithms* **14** (1999), no. 2, 153–184. MR MR1667319 (2000b:92003)
- [10] J. S. Farris, *A probability model for inferring evolutionary trees*, *Syst. Zool.* **22** (1973), no. 4, 250–256.
- [11] J. Felsenstein, *Inferring phylogenies*, Sinauer, New York, New York, 2004.
- [12] Valerie King, Li Zhang, and Yunhong Zhou, *On the complexity of distance-based evolutionary tree reconstruction*, *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms* (Baltimore, MD, 2003) (New York), ACM, 2003, pp. 444–453. MR MR1974948
- [13] Michelle R. Lacey and Joseph T. Chang, *A signal-to-noise analysis of phylogeny estimation by neighbor-joining: insufficiency of polynomial length sequences*, *Math. Biosci.* **199** (2006), no. 2, 188–215. MR MR2211625 (2007a:92048)

- [14] Michel Ledoux, *The concentration of measure phenomenon*, Mathematical Surveys and Monographs, vol. 89, American Mathematical Society, Providence, RI, 2001. MR MR1849347 (2003k:28019)
- [15] G. Liang, E. Mossel, and B. Yu, *Network topology inference through end-to-end measurements*, 2007.
- [16] E. Mossel, *Distorted metrics on trees and phylogenetic forests*, IEEE/ACM Trans. Comput. Bio. Bioinform. **4** (2007), no. 1, 108–116.
- [17] Elchanan Mossel and Sébastien Roch, *Learning nonsingular phylogenies and hidden Markov models*, Ann. Appl. Probab. **16** (2006), no. 2, 583–614. MR MR2244426
- [18] Rajeev Motwani and Prabhakar Raghavan, *Randomized algorithms*, Cambridge University Press, Cambridge, 1995. MR MR1344451 (96i:65003)
- [19] J. Ni and S. Tatikonda, *A Markov random field approach to multicast-based network inference problems*, Proceedings of the IEEE International Symposium on Information Theory, 2006, pp. 2769–2773.
- [20] ———, *Explicit link parameter estimators based on end-to-end measurements*, Forty-Fifth Annual Allerton Conference, 2007.
- [21] ———, *Network tomography based on additive metrics*, Proceedings of the 42nd Annual Conference on Information Sciences and Systems, 2008, pp. 1149–1154.
- [22] Jian Ni, Haiyong Xie, S. Tatikonda, and Y.R. Yang, *Network routing topology inference from end-to-end measurements*, INFOCOM 2008. The 27th Conference on Computer Communications. IEEE (2008), 36–40.
- [23] Francesco Lo Presti, N. G. Duffield, Joe Horowitz, and Don Towsley, *Multicast-based inference of network-internal delay distributions*, IEEE/ACM Trans. Netw. **10** (2002), no. 6, 761–775.
- [24] N. Saitou and M. Nei, *The neighbor-joining method: A new method for reconstructing phylogenetic trees*, Mol. Biol. Evol. **4** (1987), no. 4, 406–425.
- [25] Charles Semple and Mike Steel, *Phylogenetics*, Oxford Lecture Series in Mathematics and its Applications, vol. 24, Oxford University Press, Oxford, 2003. MR MR2060009 (2005g:92024)
- [26] Peter H. A. Sneath and Robert R. Sokal, *Numerical taxonomy*, W. H. Freeman and Co., San Francisco, Calif., 1973, The principles and practice of numerical classification, A Series of Books in Biology. MR MR0456594 (56 #14818)
- [27] Y. Vardi, *Network tomography: estimating source-destination traffic intensities from link data*, J. Amer. Statist. Assoc. **91** (1996), no. 433, 365–377. MR MR1394093 (97a:62050)

A Examples of Regular Delay Distributions

Below, we give two typical examples of families of distributions covered by our results. The first example is a set of continuous distributions with few parameters. The second example is a general discrete distribution. The latter is the main focus of [23].

Uniform distributions. Let $\mathcal{Q} = \{Q_\theta\}_{\theta \in \Theta}$ be the family of distributions where Q_θ is uniform on $[0, \theta]$ with $\Theta = [\underline{\theta}, \bar{\theta}]$ for some $0 < \underline{\theta} < \bar{\theta} < +\infty$. Let $\hat{w}^{(2)}$ be the estimated variance and define

$$\hat{\theta}^2 = \Psi(\hat{w}^{(2)}) = \begin{cases} \frac{\theta^2}{\bar{\theta}^2}, & \text{if } 12\hat{w}^{(2)} < \frac{\theta^2}{\bar{\theta}^2}, \\ \frac{\theta^2}{\bar{\theta}^2}, & \text{if } 12\hat{w}^{(2)} > \frac{\theta^2}{\bar{\theta}^2}, \\ 12\hat{w}^{(2)}, & \text{otherwise.} \end{cases}$$

Assume $|\hat{w}^{(2)} - w^{(2)}| \leq \delta \equiv \frac{\varepsilon \bar{\theta}^2}{12}$. From $\theta^2 - \hat{\theta}^2 = (\theta - \hat{\theta})(\theta + \hat{\theta})$, it follows easily that $|\theta - \hat{\theta}| \leq \frac{\varepsilon \bar{\theta}}{2}$. Note that

$$\|Q_\theta - Q_{\hat{\theta}}\|_1 = \int_0^{\bar{\theta}} \left| \frac{\mathbf{1}_{x \leq \theta}}{\theta} - \frac{\mathbf{1}_{x \leq \hat{\theta}}}{\hat{\theta}} \right| dx,$$

and assuming w.l.o.g. that $\theta > \hat{\theta}$ (the other case is symmetric)

$$\int_0^{\bar{\theta}} \left| \frac{\mathbf{1}_{x \leq \theta}}{\theta} - \frac{\mathbf{1}_{x \leq \hat{\theta}}}{\hat{\theta}} \right| dx = \hat{\theta} \left(\frac{1}{\hat{\theta}} - \frac{1}{\theta} \right) + (\theta - \hat{\theta}) \frac{1}{\theta} \leq 2 \frac{\theta - \hat{\theta}}{\underline{\theta}} \leq \varepsilon.$$

Therefore, \mathcal{Q} is $(\varepsilon, 2)$ -regular for any $\varepsilon > 0$.

Bounded discrete distributions. Let M be a positive integer and let $[M] = \{0, 1, \dots, M\}$. Also, let $0 < \underline{\theta} < 1$ and

$$\Theta = \left\{ \theta = (\theta_0, \theta_1, \dots, \theta_M) : 0 \leq \theta_i \leq 1, \forall i \in [M], \theta_0 > \underline{\theta}, \text{ and } \sum_{i \in [M]} i \theta_i \in [M] \right\}.$$

Denote by $\mathcal{Q} = \{Q_\theta\}_{\theta \in \Theta}$ the family of distributions on $[M]$ such that $X \sim Q_\theta$ means

$$\mathbb{P}[X = i] = \theta_i, \quad \forall i \in [M].$$

The assumption on the mean of X in the definition of Θ greatly simplifies the calculations below. It is a reasonable approximation in the standard practical case where \mathcal{Q} is a discretization of continuous densities with a large number of bins M . The assumption on θ_0 simply indicates that the distribution has been translated to “start at 0.” Define $\mu = \mathbb{E}[X]$ where $X \sim Q_\theta$ and let $\theta' = (\theta'_{-M}, \theta'_{-M+1}, \dots, \theta'_M)$ where $\theta'_{i-\mu} = \theta_i$ for all $i \in [M]$ and 0 otherwise. Note that the following holds

$$\sum_{i=-M}^M i^j \theta'_i = w^{(j)}(\theta), \quad \forall j \in [2M+1],$$

or in matrix form $\Lambda \theta' = \mathbf{w}$. From the Vandermonde structure of Λ it follows easily that $\det \Lambda \geq 1$, that is, Λ^{-1} exists, and furthermore $\|\Lambda^{-1}\|_1$ is a strictly positive constant depending on $\underline{\theta}, M$. Let $\hat{\mathbf{w}}$ be the estimate of \mathbf{w} and let $\hat{\theta}' = \Lambda^{-1} \hat{\mathbf{w}}$. Then, it follows that for any $\varepsilon > 0$ there is $\delta > 0$ such that

$$\|\hat{\theta}' - \theta'\|_1 \leq \|\Lambda^{-1}\|_1 \|\mathbf{w} - \hat{\mathbf{w}}\|_1 \leq \varepsilon,$$

whenever $\|\mathbf{w} - \hat{\mathbf{w}}\|_\infty \leq \delta$. Assume further that $\varepsilon < \underline{\theta}/2$, then we can recover an estimate $\hat{\theta}$ of θ from $\hat{\theta}'$ such that $\|\hat{\theta} - \theta\|_1 \leq \varepsilon$. Indeed, our assumptions above allow us to infer a distribution centered at 0 which we then translate to start at 0. Therefore, \mathcal{Q} is $(\varepsilon, 2M-1)$ -regular. Note that strictly speaking one should force all components of $\hat{\theta}$ to be in $[0, 1]$ and renormalize appropriately. Details are omitted.




Figure 7: Illustration of routine MINI CONTRACTOR.

B DMR Algorithm

We shall now provide an outline of the DMR algorithm. The general DMR algorithm actually allows the user to build a “forest” when the number of samples is too small. We will not use this feature here and we therefore simplify the algorithm accordingly. The input to the algorithm is a $(\tilde{\tau}, \tilde{M})$ -distorted metric \widehat{W} on n leaves. In particular, we assume that the values $\tilde{\tau}$ and \tilde{M} are known to the algorithm. We denote the true tree by $T = (V, E)$. Take $\alpha, \alpha' > 0$ and $0 < \beta, \beta' < 1$ such that

$$6 < \alpha' + 3 < \alpha < (\tilde{\alpha})^{-1},$$

and

$$(\tilde{\beta})^{-1} \tilde{M} + \tilde{\tau} < \beta \tilde{M} < \frac{1}{2} [\beta' \tilde{M} - 3\tilde{\tau}].$$

(Here it is assumed that $\tilde{M} = \omega(\tilde{\tau})$.) The details of the subroutines MINI CONTRACTOR and EXTENDER can be found in Figures 9 and 10. The reader is referred to [8] for a detailed explanation of the algorithm—which is somewhat involved. In a nutshell, for each pair of leaves u, v that are not “too far”: 1) the algorithm finds all edges sitting on the path between u and v (as illustrated in Figure 7); 2) then it derives the bipartitions corresponding to these edges by “extending” the bipartitions in a small ball around u, v (as illustrated in Figure 8).

- **Pre-Processing: Proximity Test.** Build the graph $\hat{H}_\beta = (\hat{V}_\beta, \hat{E}_\beta)$ where $\hat{V}_\beta = L$ and $(u, v) \in \hat{E}_\beta \iff \widehat{W}(u, v) < \beta \tilde{M}$;

- **Main Loop.**

– For all pairs of leaves $u, v \in \hat{V}_\beta$ such that $(u, v) \in \hat{E}_\beta$:

* **Mini Reconstruction.** Compute

$$\{\psi_j(u, v)\}_{j=1}^{r(u, v)} := \text{MINI CONTRACTOR}(\hat{H}_\beta; u, v);$$

* **Bipartition Extension.** Compute

$$\{\bar{\psi}_j(u, v)\}_{j=1}^{r(u, v)} := \text{EXTENDER}(\hat{H}_\beta, \{\psi_j(u, v)\}_{j=1}^{r(u, v)}; u, v);$$

– Deduce the tree \hat{T} from $\{\bar{\psi}_j(u, v)\}_{j=1}^{r(u, v)}$;

- **Output.** Return the resulting tree \hat{T} .




Figure 8: Illustration of routine EXTENDER.

Algorithm MINI CONTRACTOR

Input: Graph \widehat{H}_β ; Leaves u, v ;

Output: Bipartitions $\{\psi_j(u, v)\}_{j=1}^{r(u,v)}$;

- **Ball.** Let

$$\widehat{B}_{\beta'}^{(0)}(u, v) := \left\{ w \in \widehat{H}_\beta : \widehat{W}(u, w) \vee \widehat{W}(v, w) < \beta' \widetilde{M} \right\};$$

- **Intersection Points.** For all $w \in \widehat{B}_{\beta'}^{(0)}(u, v)$, estimate the point of intersection between u, v, w (distance from u), that is,

$$\widehat{\Phi}_w := \frac{1}{2} \left(\widehat{d}(u, v) + \widehat{d}(u, w) - \widehat{d}(v, w) \right);$$

- **Long Edges.** Set $S := \widehat{B}_{\beta'}^{(0)}(u, v) - \{u\}$, $x_{-1} = u$, $j := 0$;

– Until $S = \emptyset$:

- * Let $x_0 = \arg \min \{ \widehat{\Phi}_w : w \in S \}$ (break ties arbitrarily);
- * If $\widehat{\Phi}_{x_0} - \widehat{\Phi}_{x_{-1}} \geq \alpha' \widetilde{\tau}$, create a new edge by setting $\psi_{j+1}(u, v) := \{ \widehat{B}_{\beta'}^{(0)}(u, v) - S, S \}$ and let $C_{j+1} := \{x_0\}$, $j := j + 1$;
- * Else, set $C_j := C_j \cup \{x_0\}$;
- * Set $S := S - \{x_0\}$, $x_{-1} := x_0$;

- **Output.** Return the bipartitions $\{\psi_j(u, v)\}_{j=1}^{r(u,v)}$.

Figure 9: Algorithm MINI CONTRACTOR.

Algorithm EXTENDER

Input: Graph \widehat{H}_β ; Bipartitions $\{\psi_j(u, v)\}_{j=1}^{r(u, v)}$; Leaves u, v ;

Output: Bipartitions $\{\bar{\psi}_j(u, v)\}_{j=1}^{r(u, v)}$;

- For $j = 1, \dots, r(u, v)$ (unless $r(u, v) = 0$):
 - **Initialization.** Denote by $\psi_j^{(u)}(u, v)$ the vertex set containing u in the bipartition $\psi_j(u, v)$, and similarly for v ; Initialize the extended partition $\bar{\psi}_j^{(u)}(u, v) := \psi_j^{(u)}(u, v)$, $\bar{\psi}_j^{(v)}(u, v) := \psi_j^{(v)}(u, v)$;
 - **Modified Graph.** Let K be \widehat{H}_β where all edges between $\psi_j^{(u)}(u, v)$ and $\psi_j^{(v)}(u, v)$ have been removed;
 - **Extension.** For all $w \in \widehat{v}_\beta^{(i)} - (\psi_j^{(u)}(u, v) \cup \psi_j^{(v)}(u, v))$, add w to the side of the partition it is connected to in K (by definition of K , each w as above is connected to exactly one side);
- Return the bipartitions $\{\bar{\psi}_j(u, v)\}_{j=1}^{r(u, v)}$.

Figure 10: Algorithm EXTENDER.